

Mobile Augmented Reality as an Example of a Complex and Demanding Human-Centered System

Simon Julierⁱ, Steven Feinerⁱⁱ and Lawrence Rosenblumⁱⁱⁱ

Abstract— This position paper describes a recently initiated research project, the Battlefield Augmented Reality System (BARS). BARS will explore how important information can be overlaid on what users see, hear, and otherwise experience of the surrounding world as they walk through an urban environment. We argue that the development of such a human-centered system introduces many difficult challenges in a variety of areas including: accurate, wide-area tracking systems; navigation, user interface design; information filtering; wireless networking and software architectures; and societal issues for distributed collaboration.

Index terms—distributed VR; augmented reality; wearable computing; tracking and registration; interactive visualization; command and control.

1 INTRODUCTION

Recent developments in computing hardware have finally begun to make wearable AR systems feasible. With this new freedom, it becomes possible for AR systems to be used in a very wide range of applications including disaster relief, management and repair of utilities in a street and even an assistant for tourists walking through unfamiliar historical sites. However, working with large environments introduces many new challenges concerned with: navigation and wayfinding (the user knows where they are and how they can get to new desired locations); querying for information about the environment (finding out what they are looking at).

Given its ubiquitous nature, the problem of urban complexity is the subject of intense study within the field of wearable computers and a number of research avenues have been pursued. Within the military context approaches such as digital maps or “rolling compass” displays have been used [Gumm-98]. Although these methods have been demonstrated to be superior to systems based on radioed

instructions, new and better approaches are being sought. The most promising are those based on Augmented Reality (AR), where the user sees, hears, or otherwise experiences the real environment augmented by additional information.

Experimental AR prototypes have been demonstrated in task domains ranging from aircraft manufacturing [Caudell-92, Caudell-94], to image-guided surgery [Fuchs-98], and from maintenance and repair [Feiner-93, Hoff-96] to building construction [Webster-96]. These have provided many insights into how to design and use AR systems. However, mobile AR systems introduce a new set of challenging and largely unstudied research problems. First, since the user potentially operates over a very large area, conventional tracking systems that rely on bounded working volumes (e.g., [Welch-97] presents a room-sized tracking system) or a specially prepared environment [Koller-97] (the environment is populated by specially placed distinctive marks or *fiducials*) are inappropriate. Second, as a user moves through the environment their context can change dramatically depending on their position and current intent. To explore these and other issues raised by mobile augmented reality, the Naval Research Laboratory and Columbia University have embarked on a research program called the Battlefield Augmented Reality System (BARS).

BARS was motivated by the fact that with the proliferation of urbanization throughout the world, it is expected that many future military operations (such as peace keeping or hostage rescue) will occur in urban environments. These environments present many challenges. First, urban environments are extremely complicated and inherently three-dimensional. Above street level, the infrastructure of buildings may serve many different purposes (such as hospitals or communication stations) and can harbor many types of risks (such as snipers or instability due to structural damage). These features are often distributed and interleaved over several floors of a multi-floor building. Below street level, there may be a complex network of sewers, tunnels and utility systems. Cities can be confusing (especially if street signs are damaged or missing) and coordinating multiple team members can be difficult. To ensure the safety of both civilian and military personnel, it has long been argued that environmental information must

ⁱ ITT Systems and Sciences/Advanced Information Technology Division, Naval Research Laboratory. Email: julier@ait.nrl.navy.mil

ⁱⁱ Columbia University. Email: feiner@cs.columbia.edu

ⁱⁱⁱ Advanced Information Technology Division, Naval Research Laboratory. Email: rosenblu@ait.nrl.navy.mil

^{iv} Columbia University, Dept. of Computer Science, New York, NY. Email: feiner@cs.columbia.edu

be delivered to the individual user *in situ* in that environment. Despite the specific nature of this application, it touches on many areas of research including user interface design, software architecture and distributed interaction.

BARS builds from two thrusts. The first is the commercial development of wearable hardware that is finally beginning to make wearable AR systems feasible. Wearable computers are becoming small enough and powerful enough to generate 3D graphics at interactive frame rates. Head-worn displays are on their way to reaching the required resolution, brightness and form factor [Pryor-98, Spitzer-98]. Many of these components are driven by other commercial considerations (e.g., Sony's Glasstron family of head-worn displays is being developed in large part for entertainment).

The second thrust is the demonstration of a mobile AR testbed called the Touring Machine [Feiner-97, MacIntyre-99]. Developed at Columbia University, one of its experimental applications provides information about the environment to a user walking through the Columbia University campus. The user's position and orientation are continuously tracked using a real-time kinematic GPS receiver and an inertial tracker. An optical see-through head-worn display allows the user to see overlaid information about surrounding campus buildings. Using a handheld computer, the user can select buildings in the current field of view and acquire information, such as web pages of the academic department housed in the selected building.

We are extending this work by exploring how AR systems can be used to deliver detailed information about complicated, dynamic, and mutable environments. From our experiences with the Touring Machine, we conclude that significant research must be directed towards three main areas: tracking systems, the design of the information displays, and the means by which the user interacts with the system.

The structure of this paper is as follows. In Section 2 we outline the problems faced by tracking systems, and describe current solutions and research efforts. Section 3 considers user interface design issues. We examine interaction methods in Section 4 and present our conclusions in Section 5.

2 TRACKING SYSTEMS

To successfully register graphical information with the environment, an AR system must know where the user is located and what the user is looking at. However, the nature of head movements (which can be extremely rapid with peak angular accelerations of over 200deg/s^2), and the precision that is required (often sub-degree, and ideally sub-pixel level, accuracy) means that tracking is one of the most significant challenges to the development of almost any kind of responsive VR system [Azuma-97]. See-

through AR systems present even greater challenges than their VR counterparts. The fundamental reason is that the visible environment is a function of the user's *actual* position and orientation and is not calculated from any kind of tracking system. However, the graphical displays are generated using a position calculated from a tracking system. Many methods for masking latency or errors in tracking systems cannot be applied. For example, researchers have demonstrated the use of a temporal shift to delay the image of the real world to reduce the effective lag of overlaid material for a video-mixed system [Bajura-95]. However, such an approach would be highly disorienting for a user who had to interact with that world (and impossible in an optical see-through system). Accurate outdoor tracking is significantly more difficult than what is normally addressed by indoor tracking systems: tethered trackers cannot be used and fiducial-based systems are infeasible since the environment may be both unstable and unavailable for the addition of fiducials.

Given these difficulties, it has been acknowledged that any tracking solution must be a hybrid that combines several different tracking technologies [Azuma-98]. One project that is examining this problem in detail is Geospatial Registration of Information for Dismounted Soldiers (GRIDS) [GRIDS-98]. Technologies under study include the use of image-based trackers for optical flow and landmark tracking, combined with magnetometers and high precision inertial systems. To date, GRIDS has primarily focused on the problem of estimating orientation [Azuma-99, You-99].

Because the technologies for estimating position and orientation are different, the next two subsections consider the problems of estimating each of these components separately. We conclude with a discussion of estimating both parameters together.

2.1 Position

The problem of outdoor position tracking has been greatly mitigated by the development of real-time kinematic GPS systems. Current systems are capable of estimating their own positions with standard deviations of fewer than five centimeters. However, the urban environment is punishing to this type of technology. The success of a GPS system depends on its ability to observe signals from multiple satellites simultaneously. As the number of visible satellites decreases, the accuracy of the position solution declines. At least four satellites (three if the altitude is assumed to be a known constant) must be visible for GPS to yield a unique solution. However, the buildings within a city can mask a significant part of the sky and the actual number of satellites depends on the satellite constellation and local building configuration. In simple empirical trials at the Naval Research Laboratory, the number of tracked satellites fell from nine to three when the GPS receiver moved about ten meters from the middle of an open area to

the center of a five meter-wide road between three-storey research buildings. Buildings create further difficulties through multi-path effects: the signals from a satellite can be reflected from the surrounding and the road, causing multiple copies of the signal, all traversing slightly different routes, to arrive at the receiver at approximately the same time.

Some of these difficulties can be overcome by recent developments in GPS-based tracking technologies. Satellite availability can be increased by using dual-constellation GPS systems. Such systems employ both the US GPS (24 satellites) and Russian GLONASS (currently 17 satellites). Multi-path effects are being addressed in new technologies such as SnapTrack [SnapTrack-99]. SnapTrack processes measurements from a GPS receiver in conjunction with a local server on cellular telephone network. The local server provides information such as the constellation of satellites currently in view, Doppler offsets and altitude. It provides services such as sequential measurement optimization techniques to mitigate errors caused by multi-path and reflections.

Despite these advances, GPS still has a number of limitations. It will not work inside buildings, can be blocked by foliage, and is vulnerable to jamming. The additional radio signal used to provide the special differential correction information needed for real-time kinematic GPS have also been problematic since it currently uses frequency spectrum that is also allocated to other uses, including voice transmission. Since FCC regulations mandate the data to be a second-class citizen on voice bands, data transmitters must yield to avoid interference with voice. This can result in unpredictable intermittent gaps in the broadcast of differential corrections, during which accuracy plummets. The GPS industry is actively pursuing an initiative to get the FCC to allocate spectrum for RTK systems. While these issues make sole reliance on GPS problematic in production systems, we see GPS as an invaluable research tool.

Although GPS is most popular among researchers, other sensing options are feasible and are being examined. These include the use of rate sensors that integrate position and velocity sensors to estimate the movement of an individual. The Army's Land Warrior system, for example, combines a GPS receiver with a pedometer [Judd-97].

2.2 Orientation

In some respects, an accurate orientation estimate is more important than an accurate position estimate. The reason is that when targets are viewed at a long distance, the errors in registration are dominated by orientation errors. Estimating these states is even harder because orientation is mostly a function of head orientation and, as mentioned earlier, head movements can be extremely violent and unpredictable.

The most feasible solutions are those that fuse an inertial system with some kind of absolute orientation data [Foxlin-96] to mitigate gyro drift. The Intersense IS300Pro, for example, fuses a high frequency inertial loop with gravimeters (to define "vertical") and a magnetic compass. However, a recent analysis of magnetic compasses has shown that they are prone to drift [Azuma-99] in a natural outdoor environment.

2.3 Joint Approaches

Many types of tracking systems estimate position and orientation simultaneously. This is an extremely large class and includes magnetic trackers (such as the Polhemus Fastrak), ultrasound- and inertial tracking-based (the InterSense IS600Pro) and vision-based systems (e.g., the fiducial-based system described in [Koller-97]). Of these types, the latter is potentially the most useful for outdoor AR environments. There is a significant body of work on tracking naturally occurring features in an environment, e.g. [Behringer-99].

Of these different approaches, we believe a hybrid system that consists of: an RTK GPS receiver, an inertial tracking system, a natural land-mark based visual tracking system

We also note that integrating a number of disparate sensor systems together into a hybrid tracking technology introduces a number of challenges to the underlying theoretical foundations of conventional data fusion systems. Although it is rarely appreciated, the successful operation of an algorithm such as the Kalman filter [Welch-97, Foxlin-98] relies on the assumption that the unmodeled disturbances and the errors in the different sensing system are all independent of one another. However, in practice there are several problems with this assumption. First, most sensing systems perform various kinds of pre-filtering and data manipulation operations. These operations, which are often proprietary, mean that one cannot assume that errors are independent. Second, even if one were able to construct a full model of the operation of each sensor, the resulting system would be of extremely high dimension and interactive rates (updates on the order of hundreds of Hz) could not be achieved with current computing technologies. Finally, the forces that act on a system (such as a user's actual head motion) are *not* independent.

Given these difficulties, we propose that the use of robust data fusion algorithms such as Covariance Intersection (CI) [Julier-97, Uhlmann-98] need to be applied. This algorithm is similar to the Kalman filter but does not rely on the independence assumption. Using this algorithm it is possible to manipulate noises that are not independent and break up monolithic sensor fusion architectures into a set of interoperable, data fusion systems, each of which runs at its own rate.

The success of any computer-based information system is largely dependent on the quality of its user interface. Otherwise excellent information systems with cumbersome and confusing user interfaces can be of limited use. This is very important for the design of the BARS interface because the urban environment has a very high information density. A great deal of potentially interesting information (such as the names, locations, and other attributes of surrounding infrastructure) is associated with the user's immediate environment and the risk of information overload through showing too much information is high [AGARD-88]. However, almost all guidelines for the design of user interfaces for mobile AR systems are oriented towards the design of heads-up displays for aircraft and very few studies of the design for wearable systems or AR systems have been made [Billinghurst-98, AGARD-88]. One important exception was the study of a user interface for the US Army's Land Warrior (LWP) [Gumm-98].

LWP is exploring how modern computer technology can be used to provide soldiers with personal information processing systems. A wearable computer was developed with an opaque monocular display and a user interface with two options — a digital map and a "rolling compass display". The digital map is a plan view of the environment that contains icons showing the location of the user as well as other objects of interest. The rolling compass display shows a line with bearing marked on it. As the user's head turns, the line "slides" by. Icons are attached to the compass. To test the effectiveness of these interfaces, users underwent a series of test trials in a rural environment [Gumm-98]. It was found that the system improved a number of aspects of performance. However, most soldiers used the system for less than 25% of the time. Although Gumm did not identify the reasons for this pattern of use, we believe that it illustrates the need to develop good design guidelines for mobile AR systems. We have begun to explore guidelines in related fields to develop a set of hypotheses that will be studied in the BARS project. In particular, we consider the problems of environmental features, displaying routes, coordination information and conclude by discussing the issue of content management.

3.1 Environmental Information

The urban environment is extremely rich and densely packed with information. Certain tasks may require naming and highlighting buildings and other structures. When a relatively small number of buildings are to be identified, superimposing labels over the approximate center of the building is relatively simple and does not require a high performance tracking system [Feiner-97]. However, this simple strategy does not work if, for example, some buildings are partially occluded by others. Refinements such as calculating the visible part of a

building and annotating that must be considered. Furthermore, the methods for annotating small-scale features (such as the windows or doors of a building) depend on the accuracy of the tracking system. If tracking accuracy is not accurate enough that a correct feature cannot be uniquely identified, extra cues (such as instructions relative to a distinctive feature of the building) must be provided.

3.2 Routing

Routing information must show where a user is currently located and the route that must be taken to achieve a particular objective. The question of routing and the utility of different kinds of map displays have been extensively analyzed for virtual environments. For example, in [Darken-99] the relative effectiveness of a rotating map and a map with north always pointing up was considered for applications where users had to navigate through a virtual city. The conclusion was that different map displays were better suited for different types of tasks. In targeted searches, where a user was shown the location of a target of interest, rotating maps were better. This analysis suggests that dynamically updated routing information, expressed as a function of user orientation, is extremely useful. However, no such detailed studies have been carried out for AR systems that are providing continuous navigation cues.

3.3 Coordination Information

The BARS system will also present data that can be used to coordinate group activities. For example, a user may wish to know the location of colleagues with respect to themselves. As with environmental information, the means of showing this data depends on the number of users and the accuracy of the tracking systems.

3.4 Information Filtering

The amount of information that can be shown to a user in a virtual world can be overwhelming. To alleviate this problem, there is a need for an 'intelligent' filter that will determine what information is relevant to the user at a particular time. To accomplish this, the aim or goal of the user needs to be defined. For instance, in the BARS system, two goals typical of a reconnaissance mission are (i) go to a particular location; and (ii) find out as much as possible about an area. In the first situation, a route between the current and desired locations could be shown, while in the latter situation, important information about selected key objects can be displayed graphically (i.e., whether they are friendly, enemy or unknown). Some information might always be shown, such as known enemy locations and hazards (such as mines). We believe that an intelligent filter will lower the real-time update requirement of the image generator by reducing the information that it needs to display. In addition, the user

will be shown only what is relevant to them at any particular time.

4 USER INTERACTIONS

Another important research area is to investigate methods for user interaction. Within the BARS project, users are participants with the ability to make queries and send reports that can be distributed to other users and planning systems. There has been a great deal of work dedicated to the study of user interaction paradigms in VR environments (a useful taxonomy can be found in [Poupyrev-98]). However, mobile augmented reality systems present fundamentally different challenges because the interaction space is different. For example, one major challenge in VR systems is the process of constructing a user interface to drive how a user's view point will change. A range of approaches including flying, grabbing [Poupyrev-98] or using a "world in miniature" [Pausch-95] have been investigated. However, a mobile AR user can only change their position in the environment by physically moving to that new location. As explained in Subsection 3.2, the AR system can only provide prompts or cues to guide the user to the destination. Rather, the main user interactions must allow a user to query objects that can be seen in the environment and to make reports about, for example, the state of a known object in the environment or make a report about a new object. We believe that any successful interaction paradigm must span both 3D and 2D displays.

4.1 3D Interaction

Many types of systems for 3D interactions employ various types of widgets, menuing systems or button/keyboard operations. However, all of these approaches draw the user's attention away from the environment and towards manipulating physical or virtual input devices. To overcome these limitations, we propose to develop our 3D interaction system using a multi-modal input paradigm. This paradigm decomposes user interactions into a combination of several different types of input that occur concurrently. The most widely used methods combine speech with a physical gesture (such as pointing or drawing). One of the most successful promising examples of an experimental multi-modal input approach is the QuickSet system [Cohen-97]. This research has examined how a combination of speech and 2D gesture can be used to accelerate the set up process for military simulations in 2D maps. Part of the BARS process will be to see how this 2D approach can be extended to exploit the range of interactions in a full 3D space. Two key research areas remain – how is a feature of the environment selected, and how is the action to be specified?

Various types of selection methods have been proposed including virtual hand metaphors for grabbing. However, because the user is constrained to observe objects at the scale of the environment, we believe that virtual pointer

metaphors — where the user points at an object rather than grabs it — are the most appropriate for this problem domain. However, the precise question of which pointer metaphor is most appropriate is an open research problem and is likely to be a function of the task at hand. For example, when selecting pre-existing objects in the environment, selection metaphors such as an aperture [Forsberg-96], where the selection region is an adjustable cone rather than a line, might be most appropriate. However, when a user enters data about a new object the ray casting metaphor might be more appropriate because it precisely defines the user's line-of-sight. However, a line-of-sight report does not contain sufficient information to uniquely define the position of an object and other information (such as the report is about a broken window; broken windows can only occur on the side of a building; the user is pointing at one face of a particular building) must be used.

Once an object or region of the environment has been selected, the type of operation needs to be specified. Within the multi-modal paradigm, we propose to use a speech recognizer with a natural language processor.

4.2 2D Interaction

Although 3D interactions may be useful for many natural queries with the virtual environment, they are not the only interactions that might be needed. For example, if the user wishes to indicate a 2D path, it may be more effective to do so with a 2D device. Furthermore, the passive feedback of a flat surface to support the user's hand helps prevent fatigue and allows extremely accurate and precise operations to take place.

To address this, the Touring Machine provides both a see-through 3D head-worn display and the opaque 2D display of a hand-held computer. The hand-held computer's stylus and a finger-controlled trackpad on the back of the hand-held display provide input.

The question of when a 2D or a 3D display is most effective is still an open research issue and we expect that a better understanding of the interplay of these different devices and input paradigms will be one useful result of the BARS project.

5 CONCLUSIONS

This paper has introduced the problem of building wearable augmented reality systems for users operating in large, unstructured outdoor environments. There are many difficulties associated with moving AR systems from demonstrations within laboratories to systems that work outdoors.

introduced some of the issues that are being addressed in the BARS project, which explores the application of mobile AR systems to complex outdoor environments. We have provided an overview of some of the key research areas

(tracking systems, display designs and user interaction methods) where further research must be carried out.

6 ACKNOWLEDGEMENTS

This research is supported by the Office of Naval Research, Arlington, Virginia.

7 REFERENCES

- [AGARD-88] Advisory Group for Aerospace Research and Development Neuilly-Sur-Seine, "Visual Effects in the High Performance Aircraft Cockpit", *DTIC Report ADA199306*, 1998.
- [Azuma-97] R. Azuma, "A Survey of Augmented Reality", *Presence: Teleoperators and Virtual Environments*, August, Volume 4.6, pp. 355–385, 1997.
- [Azuma-98] R. Azuma, "Making Augmented Reality Work Outdoors Requires Hybrid Tracking", *Proceedings of IWARS '98 - the First International Workshop on Wearable Augmented Reality*, 1998.
- [Azuma-99] R. Azuma, B. Hoff, H. Neely III and R. Sarfaty, "A Motion-Stabilized Outdoor Augmented Reality System", *Proceedings of IEEE VR '99, Houston TX*, March, pp. 252–259, 1999.
- [Bajura-95], M. Bajura and U. Neumann, "Dynamic Registration Correction in Video-Based Augmented Reality Systems", *IEEE Computer Graphics and Applications*, Vol. 15.5, September, pp. 52–60, 1995.
- [Billinghurst-98], M. Billinghurst, J. Bowskill, N. Dyer and J. Morphet, "Spatial Information Displays on a Wearable Computer", *IEEE Computer Graphics and Applications*, Vol. 18.6, pp. 24–31, November/December, 1998.
- [Caudell-92] T. P. Caudell and D. W. Mizell, "Augmented Reality: An Application of Heads-Up Display Technology to Manual Manufacturing Processes", *Proceedings of 1992 IEEE Hawaii International Conference on Systems Sciences*, January, 1992.
- [Caudell-94] T. P. Caudell, "Introduction to Augmented Reality", *Proceedings of the Conference on Telemanipulator and Telepresence Technologies*, SPIE, Vol. 2351, 1994.
- [Cohen-97] P. Cohen, M. Johnston, D. McGee, S. Oviatt, J. Pittman, I. Smith, L. Chen and J. Clow, "QuickSet: Multimodal Interaction for Distributed Applications", *Proceedings of ACM Multimedia*, Association for Computing Machinery, New York, 1997.
- [Darken-99] R. Darken and H. Cevik, "Map Usage in Virtual Environments: Orientation Issues", *Proceedings of the 1999 IEEE VR Conference*, pp. 133–140, March, 1999.
- [Feiner-93] S. Feiner, B. MacIntyre and D. Seligmann, "Knowledge-based Augmented Reality", *Communications of the ACM*, July, Vol. 36.7, pp. 52–62, 1993.
- [Feiner-97], S. Feiner, B. MacIntyre, T. Höllerer and T. Webster, "A Touring Machine: Prototyping 3D mobile augmented reality systems for exploring the urban environment", *Proceedings of the International Symposium on Wearable Computers, Cambridge MA*, October, 1997.
- [Forsberg-96] A. Forsberg, K. Herndon, and R. Zeleznik, "Aperture-Based Selection for Immersive Virtual Environments", *Proceedings of the ACM Symposium on User Interface and Software Technology (UIST)*, pp. 95–96, 1996.
- [Foxlin-96] E. Foxlin, "Intertial Head-Tracker Sensor Fusion by a Complementary Separate-Bias Kalman Filter", *Proceedings of the VRAIS '96*, pp. 185–194, 1996.
- [Fuchs-98] H. Fuchs, M. Livingston, R. Raskar, D. Colucci, K. Keller, A. State, J. Crawford, P. Rademacher, S. Drake, and A. Meyer, "Augmented Reality Visualization for Laparoscopic Surgery", *Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI '98)*, October, 1998.
- [GRIDS-98] HRL Laboratories, UNC Chapel Hill and Raytheon Systems Company, "GRIDS: Geospatial registration of information for dismounted soldiers", <http://web-ext2.darpa.mil/ETO/wv/98overviews/index.html>, 1998.
- [Gumm-98], M. M. Gumm, W. P. Marshak, T. A. Branscome, M. Mc. Wesler, D. J. Patton and L. L. Mullins, "A Comparison of Soldier Performance Using Current Land Navigation Equipment With Information Integrated on a Helmet-Mounted Display", *ARL Report ARL-TR-1604, DTIC Report 19980527 081*, April, 1998.
- [Hix-99] D. Hix, J. E. Swan II, J. L. Gabard, M. McGee, J. Durbin and T. King, "User-Centered Design and Evaluation of a Real-Time Battlefield Visualization Virtual Environment", *IEEE Virtual Reality '99 Conference, Houston, TX*, pp. 96–103.
- [Hoff-96] W. A. Hoff, T. Lyon and K. Nguyen, "Computer Vision-Based Registration Techniques for Augmented Reality", *The Proceedings of Intelligent Robots and Control Systems XV, Intelligent Control Systems and Advanced Manufacturing*, SPIE, Vol. 2904, pp. 538–548, November, 1996.
- [Julier-97] S. Julier and J. Uhlmann, "A Non-divergent Estimation Algorithm in the Presence of Unknown Correlations", *Proceedings of the IEEE American Control Conference*, 1997.
- [Koller-97] D. Koller, G. Klinker, E. Rose, D. Breen, R. Whitaker, M. Tuceryan, "Real-time Vision-Based Camera Tracking for Augmented Reality Applications", *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST-97)*, pp. 87–94, 1997.

[MacIntyre-99] B. MacIntyre, "Exploratory Programming of Distributed Augmented Environments", *PhD thesis*, Columbia University, NY, 1999.

[Pausch-95] R. Pausch, T. Burnette, D. Brockway, M. Weiblen, "Navigation and Locomotion in Virtual Worlds via Flight into Hand-Held Miniatures", *ACM SIGGRAPH '95 Conference Proceedings, Computer Graphics*, July, 1995.

[Poupyrev-98] I. Poupyrev, S. Weghorst, M. Billinghurst and T. Ichikawa, "Egocentric Object Manipulation in Virtual Environments: Empirical Evaluation of Interaction Techniques", *Computer Graphics Forum [Eurographics Conference Issue]*, Vol. 17.3, pp. C-41-C-52, August, 1998.

[Pryor-98] H. Pryor, T. Furness and E. Viirre, E. "The Virtual Retinal Display: A New Display Technology Using Scanned Laser Light", In *Proceedings of Human Factors and Ergonomics Society, 42nd Annual Meeting*, pp. 1570-1574, 1998.

[SnapTrack-99] SnapTrack Inc., "An Introduction to SnapTrack Server-Aided GPS Technology",
<http://www.snaptrack.com/atwork.html>

[Spitzer-98] M. Spitzer, N. Rensing, P. Aquilino, M. Olson and R. McClelland, "Portable human/computer interface mounted in eyewear", *Helmet and Head-Mounted Displays III, SPIE AeroSense Conference, Vol. 3362*, 1998.

[Uhlmann-97] J. Uhlmann, S. Julier and M. Csorba, "Nondivergent Simultaneous Map Building and Localisation using Covariance Intersection", *In the Proceedings of the SPIE AeroSense Conference, Navigation and Control Technologies for Unmanned Systems II*, 1997.

[Webster-96] A. Webster, S. Feiner, B. MacIntyre, W. Massie and T. Krueger, "Augmented reality applications in architectural construction", *Designing Digital Space: An Architect's Guide to Virtual Reality*, Ed. D. Bertol, Pub. John Wiley & Sons, pp. 193–200, 1996.

[Welch-97] G. Welch, G. Bishop, "SCAAT: Incremental Tracking with Incomplete Information," *Proceedings of SIGGRAPH '97*, 1997.

[You-99] S. You, U. Neumann and R. Azuma, "Hybrid Inertial and Vision Tracking for Augmented Reality Registration", *Proceedings of IEEE VR '99, Houston TX*, March, pp. 260–267, 1999.